Hierarchical RNN with Static Sentence-Level Attention for Text-Based Speaker Change Detection

Zhao Meng, Lili Mou, Zhi Jin

Key Laboratory of High Confidence Software Technologies (Peking University), Ministry of Education, China; Software Institute, Peking University, China

zhaomeng.pku@outlook.com, doublepower.mou@gmail.com, zhijin@sei.pku.edu.cn

arXiv:1703.07713v1 [cs.CL] 22 Mar 2017

Abstract

Traditional speaker change detection in dialogues is typically based on audio input. In some scenarios, however, researchers can only obtain text, and do not have access to raw audio signals. Moreover, with the increasing need of deep semantic processing, text-based dialogue understanding is attracting more attention in the community. These raise the problem of text-based speaker change detection. In this paper, we formulate the task as a matching problem of utterances before and after a certain decision point; we propose a hierarchical recurrent neural network (RNN) with static sentence-level attention. Our model comprises three main components: a sentence encoder with a long short term memory (LSTM)-based RNN, a context encoder with another LSTM-RNN, and a static sentence-level attention mechanism, which allows rich information interaction. Experimental results show that neural networks consistently achieve better performance than feature-based approaches, and that our attention-based model significantly outperforms non-attention neural networks.[1]

Index Terms: speaker change detection, recurrent neural network, sentence-level attention

1. Introduction

Speaker change detection, or sometimes known as *speaker segmentation*, aims to find changing points of speakers in a dialogue. Specifically, a speaker change occurs when the current and the next sentences are not uttered by the same speaker [1]. Detecting speaker changes plays an important role in dialogue processing, and is a premise of speaker clustering [2], dialogue understanding [3], etc.

Traditional speaker change detection typically uses audio input [4, 5]. However, researchers may not always have access to these raw audio signals. Further, the fast development of dialogue analysis puts high demands on understanding textual data [6, 7], as human-computer conversation involves deep semantics, requiring complicated natural language processing (NLP). Vinyals et al. [6] and Li et al. [7], for example, train sequence-to-sequence neural networks to automatically generate replies in an open-domain dialogue system. They use OpenSubtitle [8] as the corpus, but assume every two consecutive sentences are uttered by different speakers (which brings much noise to their training data). Therefore, text-based speaker change detection is important to dialogue analysis, being the main focus of this paper.

Using only text to detect speaker changes raises new challenges. Previous audio-based studies depend largely on acoustic features, e.g., pitch [9] and silence points [10], which provide much information indicating the change of speakers. With textual features alone, we need deeper semantic understanding of natural language utterances.

In this paper, we formulate text-based speaker change detection as a binary sentence-pair classification problem, that is, we would like to judge whether the speaker is changing between each consecutive sentence pair (which we call a *decision point*). We also take into consideration previous and future sentences around the current decision point as context (Figure 1), serving as additional evidence.[2]

We propose a hierarchical RNN with static sentence-level attention for text-based speaker change detection. First, we use a long short term memory (LSTM)-based recurrent neural network (RNN) to capture the meaning of each sentence. Another LSTM-RNN integrates sentence information into a vector, before and after the decision point, respectively; the two vectors are combined for prediction. To better explore the context, we further apply an attention mechanism [11] over sentences to focus on relevant information during context integration. Our attention is static in that only the nearest two sentences around the decision point search for relevant information; it differs from previous research that uses dynamic attention [11, 12], which buries more important sentences under less important context. Compared with word-level attention [12, 13], our sentence-level attention is more efficient because there could be hundreds of words in the context within only a few sentences.

Our model was evaluated on transcripts of nearly 3,000 episodes of TV talk shows. We crawled data from the Cable News Network website, and extracted dialogue utterances from original html files; speaker identities are also tagged in the corpus, serving as labels. In our experiments, modern neural networks consistently outperform traditional methods that use handcrafted features. Ablation tests confirm the effectiveness of context; the proposed hierarchical RNN with sentence-level static attention can better utilize such contextual information, and significantly outperforms non-attention neural networks. The results show that our tailored model is especially suited to the task of text-based speaker change detection.

2. Related Work

Traditional speaker change detection deals with audio input and is a key step for *speaker diarization* (determining "who spoke when?") [1]. A typical approach is to compare consecutive sliding windows of input—with spectral features [14], say—using metrics including the Bayesian information criterion [15], generalized likelihood ratio [16], and Kullback-Leibler diver-

[1] We release our code (including both data pre-processing and neural models) at https://sites.google.com/site/scdrnn/

[2] Because our task is based on text (and is *not* online speaker change detection), we actually have access to the "future context" after the decision point.

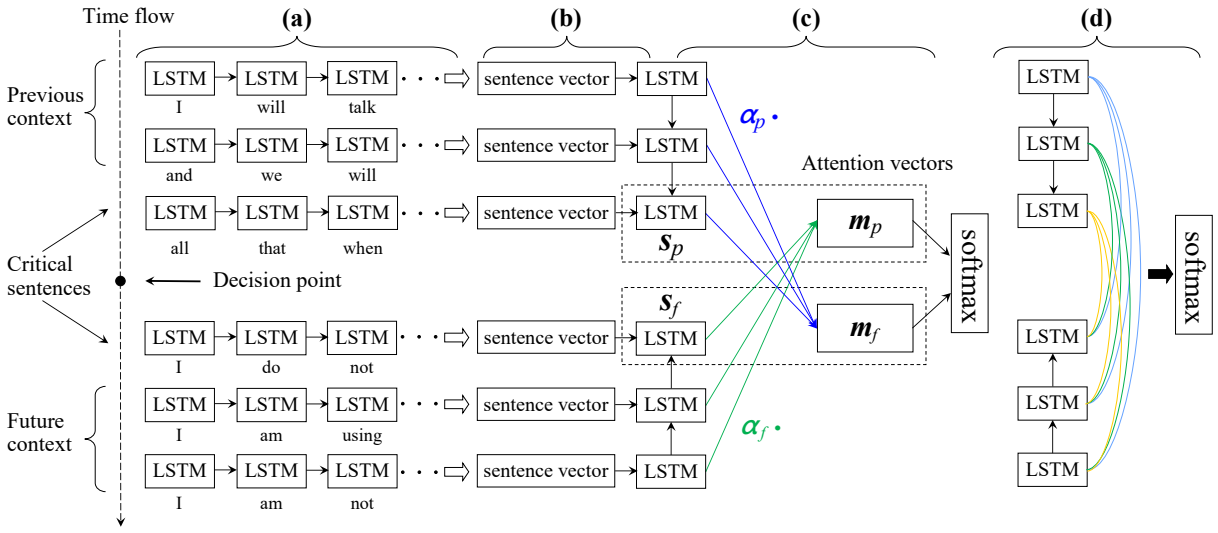

Figure 1: *The proposed neural network. (a) LSTM-RNN sentence encoder. (b) Context encoder with another LSTM-RNN. (c) Sentence-level static attention. (d) C.f. Dynamic attention. Notice that we consider text-based speaker change detection in this paper, so we have future utterances as context.*

gence [17]. Our paper differs from the above work and focuses on textual input, which is useful in various scenarios, for example, processing dialogue transcripts where speaker identities are missing (e.g., OpenSubtitle) [7], and enhancing audio speaker change detection with textual information [18].

Nowadays, text-based dialogue analysis has been increasingly important, as surface acoustic features are insufficient for semantic understanding in conversations. Previous research has addressed a variety of tasks, ranging from topic tracking [19], dialogue act classification [20], slot filling [21], to user intent modeling [22]. In our previous study, we address the problem of session segmentation in text-based human-computer conversations [23]. Without enough annotated data, we apply a heuristic matching approach; thus the task is unsupervised. Likewise, Li et al. [18] enhance audio-based speaker change detection with transcribed text, and they are also in the unsupervised regime. By contrast, this paper adopts a supervised setting as we have obtained massive, high-quality labels of speaker identities from the Cable News Network website.

As described in Section 1, we formulate our task as a sentence-pair classification problem. Previous studies utilize convolutional/recurrent neural networks (CNNs/RNNs) to detect the relationship (e.g., paraphrase and logical entailment) between two sentences [24, 25, 26]; Rocktäschel et al. [27] equip RNN with attention mechanisms. These studies do not consider contextual information. In a dialogue system, Yan et al. [28] rerank candidate replies based on a user-issued query; they enhance sentence-pair modeling with previous utterances as context. In our scenario, the context appears on both sides of the decision point, and we carefully design the neural architecture to better use such contextual information.

3. Approach

In this section, we describe the proposed approach in detail. Figure 1 illustrates our neural architecture, which consists of three parts.

- First, an LSTM-based RNN encodes each individual sentence, around the decision point, into a real-valued vector (Figure 1a and Subsection 3.1).
- Then another LSTM-based RNN integrates contextual information by going through the sentence vectors. This applies to both sides of the current decision point (Figure 1b and Subsection 3.2).
- We also add a static sentence-level attention mechanism to better utilize context by focusing on more relevant information (Figure 1c and Subsection 3.3).

3.1. Sentence Encoder

We use a recurrent neural network (RNN) with long short term memory (LSTM) units to encode a sentence as a vector (also known as a *sentence embedding*), shown in Figure 1a.

An RNN is suited for processing sequential data (e.g., a sentence consisting of several words), as it keeps a hidden state, changing at each time step based on its previous state and the current input. But vanilla RNNs with perceptron-like hidden states suffer from the problem of *vanishing or exploding gradients* [29, 30], being less effective to model long dependencies. LSTM units [31] alleviate the problem by better balancing input and its previous state with gating mechanisms.

Formally, let $\boldsymbol{x}^{(t)}$ be the embedding of the t-th word in a sentence, and $\boldsymbol{h}^{(t-1)}$ be the last step's hidden state. We have

$$\boldsymbol{i}^{(t)} = \sigma(W_i \boldsymbol{x}^{(t)} + U_i \boldsymbol{h}^{(t-1)} + \boldsymbol{b}_i) \quad (1)$$

$$\boldsymbol{f}^{(t)} = \sigma(W_f \boldsymbol{x}^{(t)} + U_f \boldsymbol{h}^{(t-1)} + \boldsymbol{b}_f) \quad (2)$$

$$\boldsymbol{o}^{(t)} = \sigma(W_o \boldsymbol{x}^{(t)} + U_o \boldsymbol{h}^{(t-1)} + \boldsymbol{b}_o) \quad (3)$$

$$\boldsymbol{g}^{(t)} = \tanh(W_g \boldsymbol{x}^{(t)} + U_g \boldsymbol{h}^{(t-1)} + \boldsymbol{b}_g) \quad (4)$$

$$\boldsymbol{c}^{(t)} = \boldsymbol{i}^{(t)} \otimes \boldsymbol{g}^{(t)} + \boldsymbol{c}^{(t-1)} \otimes \boldsymbol{f}^{(t)} \quad (5)$$

$$\boldsymbol{h}^{(t)} = \boldsymbol{o}^{(t)} \otimes \tanh(\boldsymbol{c}^{(t)}) \quad (6)$$

where W's and U's are weights, and $\boldsymbol{b}$'s are bias terms. $\otimes$ denotes element-wise product, σ the sigmoid function. $\boldsymbol{i}^{(t)}$, $\boldsymbol{f}^{(t)}$, and $\boldsymbol{o}^{(t)}$ are known as gates, and $\boldsymbol{h}^{(t)}$ is the current step's hidden state. For convenience, we use LSTM's final state (corresponding to the last word in a sentence) as the sentence embedding.

3.2. Context Encoder

Another LSTM-RNN encodes contextual information over sentence vectors, shown in Figure 1b. Since we model our task as

a matching problem, we apply the RNN to the decision point's both sides separately, the resulting vectors of which are concatenated as an input of prediction.

It should be noticed that our LSTM-RNN goes from far-away sentences to the nearest ones (called *critical sentences*) on both sides of the current decision point. We observe that nearer sentences play a more important role for prediction; that an RNN is better at keeping recent input information by its nature. Hence, our treatment is appropriate.

Besides, our neural network is hierarchical in that it composites sentences with words, and discourses[3] with sentences. It is similar to the hierarchical autoencoder in [32]. Other studies apply a single RNN over a discourse, also achieving high performance in tasks like sentiment analysis [33] and machine comprehension [34]. In our scenario, however, the sentences (either context or critical ones) are not necessarily uttered by a single speaker. Experiments in Section 4.3 show that hierarchical models are more suitable for speaker change detection.

3.3. Our Attention Mechanism

We use an attention mechanism to better utilize contextual information. Attention-based neural networks are first proposed to dynamically focus on relevant words of the source sentence in machine translation [11]. In our scenario, we would like to match a critical sentence with all utterances on the other side of the decision point (Figure 1c). That is to say, the attention mechanism is applied to the sentence level, different from other work that uses word-level attention [12, 13]. Our method is substantially more efficient because a context of several utterances could contain hundreds of words.

Considering t sentences (context size being $t-1$) before and after the decision point, respectively, we have $2t$ sentences in total, namely

$$\boldsymbol{s}_p^{(1)} \rightarrow \boldsymbol{s}_p^{(2)} \rightarrow \cdots \rightarrow \boldsymbol{s}_p^{(t)} \qquad \boldsymbol{s}_f^{(t)} \leftarrow \cdots \leftarrow \boldsymbol{s}_f^{(2)} \leftarrow \boldsymbol{s}_f^{(1)} \tag{7}$$

where subscripts p and f refer to previous and future utterances around the current decision point; the arrows indicate RNN's directions, as has been emphasized in Subsection 3.2.

In our attention mechanism, a critical sentence, e.g., $\boldsymbol{s}_p^{(t)}$, focuses on all sentences on the other side of the decision point $\boldsymbol{s}_f^{(1)}, \cdots, \boldsymbol{s}_f^{(t)}$, and aggregates information weighted by a probabilistic distribution $\boldsymbol{\alpha}_p \in \mathbb{R}^t$, i.e.,

$$\tilde{\alpha}_p^{(i)} = \boldsymbol{u}^T \tanh\left(W[\boldsymbol{s}_p^{(t)}; \boldsymbol{s}_f^{(i)}]\right) \tag{8}$$

$$\alpha_p^{(i)} = \mathrm{softmax}\left(\tilde{\alpha}_p^{(i)}\right) = \frac{\exp\{\tilde{\alpha}_p^{(i)}\}}{\sum_{j=1}^{t} \exp\{\tilde{\alpha}_p^{(j)}\}} \tag{9}$$

Here, $\boldsymbol{s}_p^{(t)}$ is concatenated with $\boldsymbol{s}_f^{(i)}$, processed by a two-layer perceptron (with parameters W and $\boldsymbol{u}$). $\tilde{\alpha}_i$ is a real-valued measure, normalized by the $\mathrm{softmax}$ activation function to give the probability α_i. The aggregated information, known as an *attention vector*, is

$$\boldsymbol{m}_p = \sum_{i=1}^{t} \alpha_p^{(i)} \cdot \boldsymbol{s}_f^{(i)} \tag{10}$$

Likewise, the other critical sentence $\boldsymbol{s}_f^{(t)}$ yields an attention vector $\boldsymbol{m}_f$. They are concatenated along with LSTM's output of the critical sentences for prediction. In other words, the input of $\mathrm{softmax}$ is $[\boldsymbol{s}_p^{(t)}; \boldsymbol{s}_f^{(t)}; \boldsymbol{m}_p; \boldsymbol{m}_f]$.

[3] A *discourse* can be viewed as a sequence of sentences.

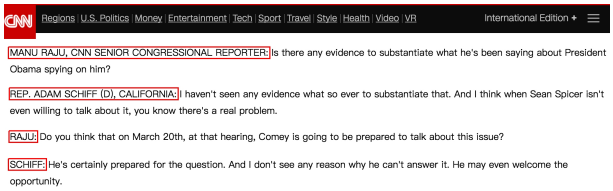
CNN
Regions | U.S. Politics | Money | Entertainment | Tech | Sport | Travel | Style | Health | Video | VR
International Edition +
MANU RAJU, CNN SENIOR CONGRESSIONAL REPORTER: Is there any evidence to substantiate what he's been saying about President Obama spying on him?
REP. ADAM SCHIFF (D), CALIFORNIA: I haven't seen any evidence what so ever to substantiate that. And I think when Sean Spicer isn't even willing to talk about it, you know there's a real problem.
RAJU: Do you think that on March 20th, at that hearing, Comey is going to be prepared to talk about this issue?
SCHIFF: He's certainly prepared for the question. And I don't see any reason why he can't answer it. He may even welcome the opportunity.

Figure 2: *An example of the Cable News Network website. Red rectangles highlight speaker identities, which serve as labels.*

	Train	Validation	Test	Total
#episodes	2,284	285	287	2,856
#utterances	1,202,789	152,068	148,939	1,503,796

Table 1: *Statistics of our corpus.*

It should be pointed out that, our attention is static, as only critical sentences search for relevant information using Equations (8)–(10). It resembles a variant in [27], but differs from common attention where two LSTMs interact dynamically along their propagations (Figure 1d). Such approach, however, may bury the critical sentences; it leads to slight performance degradation in our experiment, as we shall see in Section 4.3. By contrast, our model puts more emphasis on the critical sentences.

4. Experiments

4.1. Dataset Collection

In our experiments, we crawled transcripts of nearly 3,000 TV talk shows from the Cable News Network website[4] (example shown in Figure 2), and extracted main content from the original html files. We used the Natural Language Toolkit[5] to segment sentences and words. The transcripts contain speaker identities (highlighted by red rectangles), with which we induced speaker changes in the dialogue. Preprocessing code is released on our website (Footnote 1).

The crawled dataset comprises 1.5M utterances. We split training, validation, and test sets by episodes (TV shows) at a ratio of 8:1:1. In other words, each episode appears in either the training set, or the val/test sets. This prevents utterance overlapping between training and prediction, and thus is a more realistic setting than splitting by utterances. Table 1 presents relevant statistics.

We notice that, our corpus is larger than previous ones by magnitudes: the 1997 HUB4 dataset, for example, is 97 hours long, whereas our TV shows are estimated to be 3,000 hours,[6] which are more suitable for training deep neural networks.

In the dataset, speaker changes count to 25%. We thus used F_1-measure in addition to accuracy as metrics. Here, $F_1 = 2P \cdot R/(P+R)$, where $P = \frac{\#\text{correctly detected changes}}{\#\text{detected changes}}$ is the *precision*, and $R = \frac{\#\text{correctly detected changes}}{\#\text{all changes}}$ is the *recall*.

4.2. Settings

We set all neural layers, including word embeddings, to 200 dimensional. Since our dataset is large, we randomly initialized word embeddings, which were tuned during training. We used the Adam optimizer [35] with mini-batch update (batch size being 100). Other hyperparameters were chosen by validation: dropout rate from $\{0.1, 0.3\}$ and initial learning rate from

[4] `http://transcripts.cnn.com`
[5] `http://www.nltk.org`
[6] Each episode roughly lasts for an hour.

Model	Acc.	F_1	P	R
Random guess	61.8	25.4	26.0	25.0
Logistic regression w/ (uni+bi)-gram	80.5	50.9	73.0	39.0
DNN w/ (uni+bi)-gram	76.6	56.5	54.4	58.8
CNN w/o context	77.8	57.8	56.8	58.9
RNN w/o context	83.3	63.9	72.5	57.1
RNN w/ context (non-hierarchical)	83.7	65.7	72.6	60.0
RNN w/ context (hierarchical)	85.1	69.2	74.6	64.6
+ static attention	**89.2**	**78.4**	**81.5**	**75.6**

Table 2: *Model performance (in percentage). Here, LSTM units are used in RNN, but omitted in the table for brevity. The context size is 8, chosen by validation (deferred to Figure 3).*

$\{3 \times 10^{-4}, 9 \times 10^{-4}\}$.

We also had several baselines with handcrafted features: we extracted unigram and bigram features of the critical sentences as two vectors, which are concatenated for prediction. We applied logistic regression and a 3-layer deep neural network (DNN) as the classifier; the former is a linear model whereas the latter is nonlinear. DNN's hidden dimension was set to 200, which is the same as our neural network.

A convolutional neural network (CNN) is also included for comparison. It adopts a window size of 3, and a max-pooling layer aggregates extracted features. All competing neural models (including DNN and CNN) were tuned in the same manner, so our comparison is fair.

4.3. Performance

Table 2 presents the performance of our model as well as baselines. As shown, all modern neural networks (CNNs/RNNs with word embeddings) are consistently better than methods using handcrafted features of unigrams and bigrams. Because we have applied a 3-layer DNN to these features, we believe the performance improvement is not merely caused by using a better classifier, but the automatic feature/representation learning nature of modern neural networks.

For neural network-based sentence encoders, we compared LSTM-RNN with CNN.[7] Results show that LSTM-RNN outperforms CNN by 6% in terms of both accuracy and F_1-measure.

To cope with context, the simplest approach, perhaps, is to use a word-level RNN as above to go through surrounding utterances of the critical sentences, denoted as *non-hierarchical* in Table 2. Using contextual information yields an F_1 improvement of 2%. This controlled experiment validates the usefulness of context in speaker change detection. The hierarchical RNN introduced in Sections 3.1 and 3.2 further improves the F_1 measure by 3%. With sentence-level static attention, our model achieves the highest performance of 89.2% accuracy and 78.4% F_1-measure.

We would like to have in-depth analysis regarding how the context size and different attention mechanisms affect our model. The context size was chosen by validation from $\{1, 2, 4, 8\}$.[8] As shown in Figure 3, even a single context sentence (on each side of the decision point) improves the performance by 2%; with more surrounding utterances, the performance grows gradually. Moreover, attention-based neural networks significantly outperform non-attention models by a margin of 10%. We also tried a dynamic sentence-by-sentence attention mechanism, similar to most existing work [11, 12, 13].

[7] *CNN* here should not be confused with the Cable News Network.

[8] Due to efficiency concerns, we did not try larger context sizes.

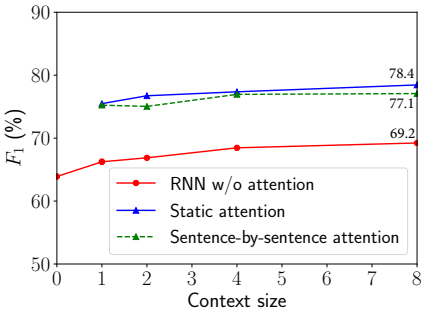

Figure 3: *Effect of the context size and different attention mechanisms. The context size is the number of utterances, excluding the critical sentence, on each side of the decision point; a size of 0 refers to no context.*

ID	Utterances	Speaker Changes? Predicted	Truth	Correct?
1	there 's no question the deficit halts on both sides of the aisles .			
2	cbo , wall street , everyone will have a say into this , including workers and future retirees .	No	No	✓
3	and your saying it 's both parties ?	Yes	Yes	✓
4	absolutely .	Yes	Yes	✓
5	the key , though , is the first six months .	No	Yes	✗
6	you say it 's not going to get through in these first couple months .	No	No	✓

Table 3: *Case study. Colors indicate speaker identities (they are used to infer groundtruth speaker changes, but cannot be seen during prediction).*

As analyzed in Section 3.3, such model buries critical sentences and thus slightly hurts the performance by 1–2% F_1-measure (green dashed line in Figure 3). The experiments verify the effectiveness of our hierarchical RNN with sentence-level, static attention.

Case study. Table 3 showcases a dialogue snippet. In the example, our model makes an error as it fails to detect the change between Sentences 4–5. However, it is even hard for humans to judge this particular change, because the word *absolutely* also goes fluently into the next speaker. For other utterances with more substance, the neural network correctly joins Sentences 1–2 and 5–6, as well as segments Sentences 2–3 and 3–4, showing that our proposed model can effectively capture the semantics of these sentences.

5. Conclusion

In this paper, we proposed a static sentence-level attention LSTM-RNN for text-based speaker change detection. Our model uses an LSTM-RNN to encode each utterance into a vector, based on which another LSTM-RNN integrates contextual information, before and after a particular decision point, respectively. A static sentence-level attention mechanism is also applied to enhance information interaction.

We crawled dialogue transcripts from Cable News Network TV talk shows for evaluation. Experimental results demonstrate the effectiveness of our approach. In particular, in-depth analysis validates that contextual information is indeed helpful for speaker change detection, and that our tailored model can make better use of context than other neural networks.

6. References

[1] X. Anguera Miro, S. Bozonnet, N. Evans, C. Fredouille, G. Friedland, and O. Vinyals, "Speaker diarization: A review of recent research," *IEEE Transactions on Audio Speech and Language Processing*, vol. 20, no. 2, pp. 356–370, 2012.

[2] R. Sinha, S. E. Tranter, M. J. Gales, and P. C. Woodland, "The Cambridge University March 2005 speaker diarisation system," in *Proceedings of INTERSPEECH*, 2005, pp. 2437–2440.

[3] N. Mrkšić, D. Ó Séaghdha, B. Thomson, M. Gasic, P.-H. Su, D. Vandyke, T.-H. Wen, and S. Young, "Multi-domain dialog state tracking using recurrent neural networks," in *Proceedings of the Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, vol. 2, 2015, pp. 794–799.

[4] S. Chen and P. Gopalakrishnan, "Speaker, environment and channel change detection and clustering via the Bayesian information criterion," in *Proceedings of DARPA Broadcast News Transcription and Understanding Workshop*, 1998, pp. 127–132.

[5] S. Kwon and S. S. Narayanan, "Speaker change detection using a new weighted distance measure," in *Proceedings of INTERSPEECH*, 2002, pp. 2537–2540.

[6] O. Vinyals and Q. Le, "A neural conversational model," in *ICML Workshop*, 2015.

[7] J. Li, M. Galley, C. Brockett, J. Gao, and B. Dolan, "A diversity-promoting objective function for neural conversation models," in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2016, pp. 110–119.

[8] J. Tiedemann, "News from OPUS-A collection of multilingual parallel corpora with tools and interfaces," in *Proceedings of Recent Advances in Natural Language Processing*, 2009, pp. 237–248.

[9] L. Lu and H.-J. Zhang, "Unsupervised speaker segmentation and tracking in real-time audio content analysis," *Multimedia Systems*, vol. 10, no. 4, pp. 332–343, 2005.

[10] T. Kemp, M. Schmidt, M. Westphal, and A. Waibel, "Strategies for automatic segmentation of audio data," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 2000, pp. 1423–1426.

[11] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *Proceedings of the International Conference on Learning Representations*, 2015.

[12] T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2015, pp. 1412–1421.

[13] B. Liu and I. Lane, "Attention-based recurrent neural network models for joint intent detection and slot filling," in *Proceedings of INTERSPEECH*, 2016, pp. 685–689.

[14] L. Lu and H.-J. Zhang, "Speaker change detection and tracking in real-time news broadcasting analysis," in *Proceedings of the Tenth ACM International Conference on Multimedia*, 2002, pp. 602–610.

[15] K. Mori and S. Nakagawa, "Speaker change detection and speaker clustering using VQ distortion for broadcast news speech recognition," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 2001, pp. 413–416.

[16] H. Gish, M.-H. Siu, and R. Rohlicek, "Segregation of speakers for speech recognition and speaker identification," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 1991, pp. 873–876.

[17] M. A. Siegler, U. Jain, B. Raj, and R. M. Stern, "Automatic segmentation, classification and clustering of broadcast news audio," in *Proceedings of DARPA Speech Recognition Workshop*, 1997.

[18] R. Li, T. Schultz, and Q. Jin, "Improving speaker segmentation via speaker identification and text segmentation." in *Proceedings of INTERSPEECH*, 2009, pp. 904–907.

[19] S. Kim, R. E. Banchs, and H. Li, "Wikipedia-based kernels for dialogue topic tracking," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, 2014, pp. 131–135.

[20] S.-S. Shen and H.-Y. Lee, "Neural attention models for sequence classification: Analysis and application to key term extraction and dialogue act detection," in *Proceedings of INTERSPEECH*, 2016, pp. 2716–2720.

[21] N. T. Vu, "Sequential convolutional neural networks for slot filling in spoken language understanding," in *Proceedings of INTERSPEECH*, 2016, pp. 3250–3254.

[22] Y.-N. Chen, M. Sun, A. I. Rudnicky, and A. Gershman, "Unsupervised user intent modeling by feature-enriched matrix factorization," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 2016, pp. 6150–6154.

[23] Y. Song, L. Mou, R. Yan, L. Yi, Z. Zhu, X. Hu, and M. Zhang, "Dialogue session segmentation by embedding-enhanced texttiling," in *Proceedings of INTERSPEECH*, 2016, pp. 2706–2710.

[24] B. Hu, Z. Lu, H. Li, and Q. Chen, "Convolutional neural network architectures for matching natural language sentences," in *Advances in Neural Information Processing Systems*, 2014, pp. 2042–2050.

[25] L. Mou, R. Men, G. Li, Y. Xu, L. Zhang, R. Yan, and Z. Jin, "Natural language inference by tree-based convolution and heuristic matching," in *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, vol. 2, 2016, pp. 130–136.

[26] S. R. Bowman, G. Angeli, C. Potts, and C. D. Manning, "A large annotated corpus for learning natural language inference," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2015, pp. 632–642.

[27] T. Rocktäschel, E. Grefenstette, K. M. Hermann, T. Kočiskỳ, and P. Blunsom, "Reasoning about entailment with neural attention," in *Proceedings of the International Conference on Learning Representations*, 2016.

[28] R. Yan, Y. Song, and H. Wu, "Learning to respond with deep neural networks for retrieval-based human-computer conversation system," in *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2016, pp. 55–64.

[29] Y. Bengio, P. Simard, and P. Frasconi, "Learning long-term dependencies with gradient descent is difficult," *IEEE Transactions on Neural Networks*, vol. 5, no. 2, pp. 157–166, 1994.

[30] S. Hochreiter, "The vanishing gradient problem during learning recurrent neural nets and problem solutions," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 6, no. 02, pp. 107–116, 1998.

[31] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[32] J. Li, T. Luong, and D. Jurafsky, "A hierarchical neural autoencoder for paragraphs and documents," in *Proceedings of the Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, 2015, pp. 1106–1115.

[33] J. Xu, D. Chen, X. Qiu, and X. Huang, "Cached long short-term memory neural networks for document-level sentiment classification," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2016, pp. 1660–1669.

[34] M. Seo, A. Kembhavi, A. Farhadi, and H. Hajishirzi, "Bidirectional attention flow for machine comprehension," *arXiv preprint arXiv:1611.01603*, 2016.

[35] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proceedings of the International Conference for Learning Representations*, 2014.